



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Roland M. HOCHMUTH et al. Docket No.: 10017761-1  
Serial No.: 10/001,430 Examiner: Ryan R. Yang  
Filing Date: October 31, 2001 Art Unit: 2628  
Title: System and Method for Communicating Graphics Image Data over a  
Communication Network

DECLARATION OF ROBERT P. MARTIN  
UNDER 37 C.F.R. SECTION 1.131

Mail Stop: **Amendment**  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

I, Robert P. Martin, a co-inventor of the above-referenced patent application, enclose hereto as Exhibit A a true copy of an invention disclosure form, with dates and portions redacted as noted, which was received by the Legal-Intellectual Property department of Hewlett-Packard Company on a date prior to September 21, 2001, in the ordinary course of business as part of Hewlett-Packard Company's invention disclosure program, and which indicates a conception of the invention which is the subject of the above-referenced patent application on a date(s) prior to September 21, 2001.

The invention that is the subject matter of the above-referenced application was conceived and reduced to practice while working in the United States for Hewlett-Packard Company. The invention disclosure form attached hereto as Exhibit A includes an explanation of the subject matter of the claims of the present application. Specifically, at least pages 3-8 of the invention disclosure form include an explanation of the subject matter of the claims of the present application.

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as First Class mail in an envelope addressed to: Mail stop: Amendment, Commissioner for Patents, P.O. 1450, Alexandria, V.A. 22313-1450.

Name of Person Signing Certificate: Cindy C. Dioso

Signature: Cindy C. Dioso Date of Signing: 4/20/2007

U.S. Patent Application Serial No. 10/001,430

PATENT

On or about September 27, 2001, I received a draft of the above-referenced patent application which I reviewed so that any comments or changes I had regarding the draft patent application could be incorporated into the final draft of the above-referenced patent application.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

## SIGNATURE

Robert P. Martin

Signature



Date:

April 18, 2007

Citizenship:

United States of America

Residence:

Fort Collins, Colorado

**EXHIBIT A**



## INVENTION DISCLOSURE

TOD - SC

PAGE ONE OF

PDNO

10017761

DATE RCVD

Redacted

ATTORNEY

KmH

Instructions: The information contained in this document is **COMPANY CONFIDENTIAL** and may not be disclosed to others without prior authorization. Submit this disclosure to the HP Legal Department as soon as possible. No patent protection is possible until a patent application is authorized, prepared, and submitted to the Government.

## Descriptive Title of Invention:

Network Attached Graphics Adapter

## Name of Project:

Scalable Computing Systems

## Product Name or Number:

IVA (Internet Visualization Architecture)

Was a description of the invention published, or are you planning to publish? If so, the date(s) and publication(s):

Redacted

Was a product including the invention announced, offered for sale, sold, or is such activity proposed? If so, the date(s) and location(s):

Redacted

Was the invention disclosed to anyone outside of HP, or will such disclosure occur? If so, the date(s) and name(s):

Redacted

If any of the above situations will occur within 3 months, call your IP attorney or the Legal Department now at 1-898-4919 or 970-898-4919.

Was the invention described in a lab book or other record? If so, please identify (lab book #, etc.):

Was the invention built or tested? If so, the date:

Redacted

Was this invention made under a government contract? If so, the agency and contract number:

Redacted

Description of Invention: Please preserve all records of the invention and attach additional pages for the following. Each additional page should be signed and dated by the inventor(s) and witness(es).

- A. Description of the construction and operation of the invention (include appropriate schematic, block, & timing diagrams; drawings; samples; graphs; flowcharts; computer listings; test results; etc.)
- B. Advantages of the invention over what has been done before.
- C. Problems solved by the invention.
- D. Prior solutions and their disadvantages (if available, attach copies of product literature, technical articles, patents, etc.).

Signature of Inventor(s): Pursuant to my (our) employment agreement, I (we) submit this disclosure on this date: [ ]

|              |                    |           |                 |                   |
|--------------|--------------------|-----------|-----------------|-------------------|
| Redacted     | Roland Hochmuth    |           | Redacted        | Redacted          |
| Employee No. | Name               | Signature | Telnet Mailstop | Entity & Lab Name |
| Redacted     | John Marks         |           | Redacted        | Redacted          |
| Employee No. | Name               | Signature | Telnet Mailstop | Entity & Lab Name |
| Redacted     | Robert Paul Martin |           | Redacted        | Redacted          |
| Employee No. | Name               | Signature | Telnet Mailstop | Entity & Lab Name |
| Employee No. | Name               | Signature | Telnet Mailstop | Entity & Lab Name |

(If more than four inventors, include additional information on another copy of this form and attach to this document)

Redacted

J.C.

PAGE \_\_\_\_\_ OF \_\_\_\_\_

The invention was first explained to, and understood by, me (us) on this date: [ ]

**Date of Signature**

## Zip

**Redacted**

Zip

**Redacted**

**Redacted**

Zip

**Redacted**

Zip

## Citizenship

2

# Description

This invention disclosure describes a graphics adapter with frame-buffer compression engine and a network interface. A network attached graphics adapter is a device similar to a traditional graphics adapter that adds to or replaces the standard video output port or ports, such as a DVI port, with a network interface port or ports such as an Ethernet or Infiniband port. The network interface port can be used for transmitting the contents of the frame-buffer in either compressed or uncompressed form or any region of the graphics adapter's memory to one or more remote network attached devices such as a compute node with a network interface card. This capability allows multiple users to remotely access and view graphics content created by graphics accelerators in compute nodes, such as **high-availability (HA) servers**, located anywhere, such as in a "**back room**", "**glass room**" or an "**Internet Data Center (IDC)**", using existing **network** infrastructure and display devices, such as "**thin clients**" or "**appliances**".

Current display refresh on the graphics adapter works as follows:

While (TRUE)

    If start of vertical retrace

        For all pixels in frame-buffer

            Read pixel from frame-buffer

            Convert pixel to output format

            Transmit out analog/digital port

The functionality listed above would be replaced with the following:

While (TRUE)

    If new update is requested a time threshold has been reached

        If there has been a change in some portion of the frame-buffer

            If compression of the frame-buffer enabled

                Compress the frame-buffer

                Convert the compressed buffer to Internet Protocol (IP) or other network protocol

                Transmit the compressed frame-buffer out the network interface.

        Else

            Convert the frame-buffer to IP or other network protocol

            Transmit the uncompressed frame-buffer out the network interface port.

The algorithm listed above could be combined with algorithms presented in the invention disclosure titled "Algorithms for removing redundant data in an Internet Visualization Appliance" which involved imposing a tiling on the image and keeping track of the tiles that have been modified.

The algorithm listed above could also be modified slightly to handle capabilities such as overlay planes and window id planes and cursor. In this scenario, multiple regions of the frame-buffer would be simultaneously read, for example the primary display surface and overlay plane. From this data the displayed image would be assembled and compressed by the frame-buffer compression engine. The following algorithm demonstrates this functionality.

While (TRUE)

    If new update is requested a time threshold has been reached

        If there has been a change in some portion of the primary display surface or overlay planes

If compression is enabled

- Coalesce and compress the buffer resulting from the primary and overlay planes  
And cursor location.

- Convert the compressed buffer to Internet Protocol (IP) or other network protocol
- Transmit the compressed frame-buffer out the network interface.

Else

- Convert the coalesced buffer to IP or other network protocol
- Transmit the uncompressed buffer out the network interface port.

The operation of coalescing involves reading the pixel from the higher priority buffer or determining the buffer that has precedence and performing color LUT conversions and also merging in the cursor.

One of the disadvantages of current display refresh technology is that the data that is sent out the frame-buffer is not tagged to allow delivery of the frame-buffer contents to multiple destinations. Consequently, graphics adapters are usually not used in "server" like environments where multiple users share a single machine. As a result, technologies like Citrix Metaframe are used that send the graphics protocol over the network which must be rendered by the client.

There have been graphics cards created with multiple video output ports. With these types of video cards it might be possible to have one user using one region of the frame-buffer and another user using another region of the frame-buffer. It may also be possible to load a computer with more than one graphics adapters and have a one-to-one mapping between users and graphics adapters. However, the number of physical ports on the graphics adapter and number of graphics cards in the system limits these approaches.

Using this invention it becomes possible for multiple users/clients to share or time-slice a graphics card and have their own unique session and/or desktop. Given a graphics adapter with it's own video memory associate a unique region of the graphics adapter memory with each unique network device. For example, if a graphics adapter has 32 Mbytes of memory the first 4 Mbytes might be assigned to network device 1, the second 4 Mbytes may be assigned to network device 2, and so on. The algorithm listed below would then be used to update each network-attached device:

While (TRUE)

- If new update is requested by a specific network device or a time threshold has been reached

- If the frame-buffer for the specific network device has been modified

- Configure the frame-buffer compression engine to compress the frame-buffer  
for the network device.

- Compress the frame-buffer for the network device.

- Convert to IP

- Configure the network interface unit to send the compressed frame-buffer to the  
network device.

- Transmit the compressed frame-buffer to the network device out the network  
Interface port.

Another possible use of this invention is in the area of collaboration where multiple users share a common desktop or window within a frame-buffer. In this scenario the same frame-buffer or same region of the frame-buffer would be compressed and transmitted to multiple network-attached devices.

# Advantages over Prior Art

- Off-loads the compression, read back, and transmittal of the frame-buffer from the host system or CPU. Currently, to transfer an image rendered in a graphics adapter to a remote network attached device the frame-buffer would be required to be read back from the graphics card into system memory, compressed, and sent over the host system's network interface. Alternatively, a PCI board or appliance could be used as identified in the invention disclosure called "Internet Graphics Appliance". However, this invention has additional benefits beyond the IGA approach, which are listed below.
- Transmittal of the contents of the frame-buffer only occurs when the contents of the frame-buffer change. Current graphics adapters transmit the entire contents of the frame-buffer to the monitor at a continuous refresh rate independent of whether or not any pixels within the frame-buffer have changed in value. This is necessary because monitors and other display devices do not have their own local display refresh memory. As a result, the overhead of display refresh imposes a constant drain on the graphics adapters memory sub-system, which significantly impacts graphics rendering performance since the same memory is being accessed by the rendering engine and display refresh unit. This invention decouples the actual display device refresh from the graphics adapter which could free up significant memory bandwidth when actual frame-buffer image updates are occurring less often than the video refresh rate of the display device. It is noted that if a network fabric such as Ethernet were used for display refresh, the capability presented here would be dependent on the remote network attached monitor or display device to be able to persist the transmitted frame-buffer in its own local memory and to refresh from a local memory source. Currently this technology does not exist and this is the subject of another invention disclosure that is in the process of being submitted.
- Transmittal of the contents of the frame-buffer can be limited to regions of interest. Current video output ports on graphics card transmit the entire contents of the frame-buffer at a constant monitor refresh rate such as 75 Hz independent of the portions of the frame-buffer that have been modified. Since the graphics adapter renders the pixels into the frame-buffer it should be able to easily keep track of the extents of the changes or tiles or regions that have changed within the frame-buffer very easily. The graphics adapter needs to only transmit the regions that have changed.
- Transmittal of frame-buffer contents or some other region of the graphics adapters local memory could be tagged/addressed and sent to more than one network attached device. This would allow multiple users of a graphics card to be supported simultaneously with their own session or in a collaborative environment. Current video output ports on graphics cards transmit the contents of the frame-buffer out the video port to whatever is plugged into it. The contents of the frame-buffer are not tagged and transmitted to a destination. This makes it very difficult to support multiple users or multiple display devices simultaneously using the same graphics adapter.
- Lower system cost and higher efficiency can be obtained in comparison to the invention disclosure "Internet Graphics Appliance" (IGA) or PCI board solution. The IGA is a separate appliance that grabs the incoming frame-buffer at refresh rates and performs many operations on the incoming pixel stream to identify if the current frame has changed and the regions or tiles that have changed. These changes could be more easily and efficiently tracked by a graphics adapter and could be used to identify when data in the frame-buffer needs to be transmitted and the regions that are required to be transmitted. Additionally, a separate device with it's own frame-buffer memory and electrical components would not be required which would further reduce the overall system cost. It is also noted that this capability does not currently exist within a graphics adapter. The invention disclosure titled "Internet Graphics Appliance" presents a methodology that would allow the frame-buffer or region of the frame-buffer to be set over the network to a remote display based on technologies and approaches currently available. The IGA invention could be viewed as a technology that could be used while transitioning from current technologies to the approach identified in this invention.



- Graphics adapter frame-buffer can be viewed remotely anytime, anyplace, and anywhere with no distance limitations. Current video out technologies has distance limitations and cannot be viewed from any network-attached device.
- Graphics Adapter with Frame-Buffer Compression Engine and Network Interface.

## Possible Embodiments

This section describes possible embodiments of the invention:

- The frame-buffer compression engine and networking functionality may be implemented in a separate chip or chips on the graphics adapter as shown in Figure 1 Variation 1. The frame-buffer compression engine and networking functionality may be implemented in a custom ASIC chip, DSP, FPGA, or other alternative. It may be possible to easily design a graphics board with this type of functionality on it without modifying the currently available graphics chips.
- The frame-buffer compression engine and networking functionality may be implemented as a block or blocks within the graphics chip as shown in Figure 2 Variation 2. In this design scenario current graphics chips would have to be modified significantly to add this type of functionality.
- The video and DVI port may be eliminated from the graphics adapter completely as shown in Figure 3 Variation 3. This would reduce the cost of the graphics adapter.
- The PCI or AGP interface may be eliminated from the graphics adapter in which case the graphics commands themselves would be received from the network interface as shown in Figure 4 Variation 4. However, the network connection would have to be significantly faster than what is currently available.
- The PCI or AGP interface could be substituted with another type of network interface such as Infiniband or Fiber Channel or another type of network technology that may become available. This embodiment is not shown in a diagram, but is easily created by substituting in the diagrams everywhere where is states PCI or AGP.
- Graphics Adapter with Frame-Buffer Compression Engine Only: In this embodiment the graphics adapter only has the frame-buffer compression engine and does not have a network interface. In this embodiment the graphics card compresses the frame-buffer. The compressed frame-buffer would be sent back to the host system and transmitted over the network by the host to the client.

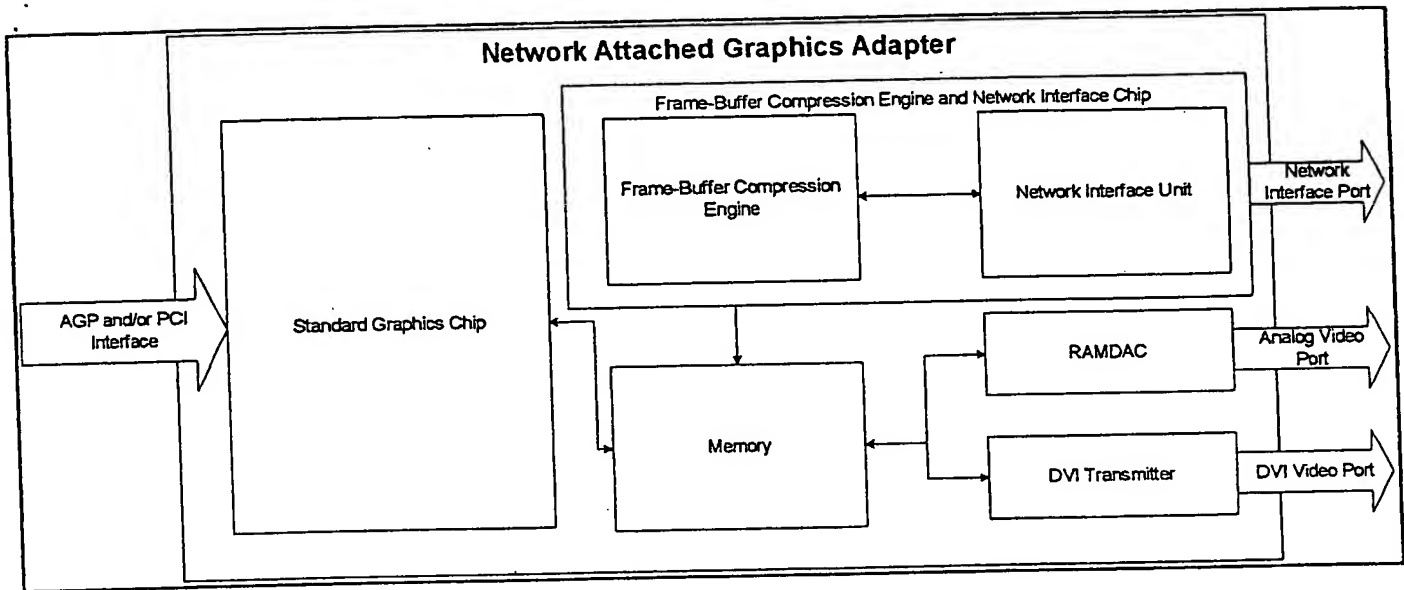


Figure 1 Variation 1

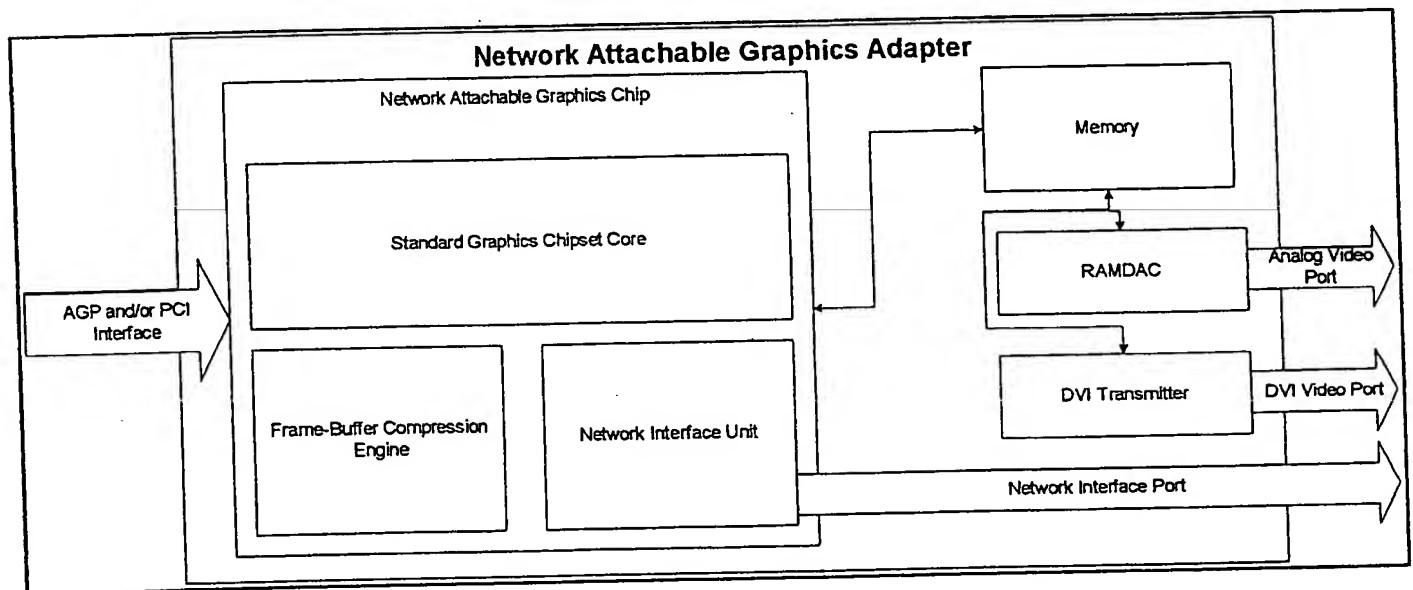


Figure 2 Variation 2

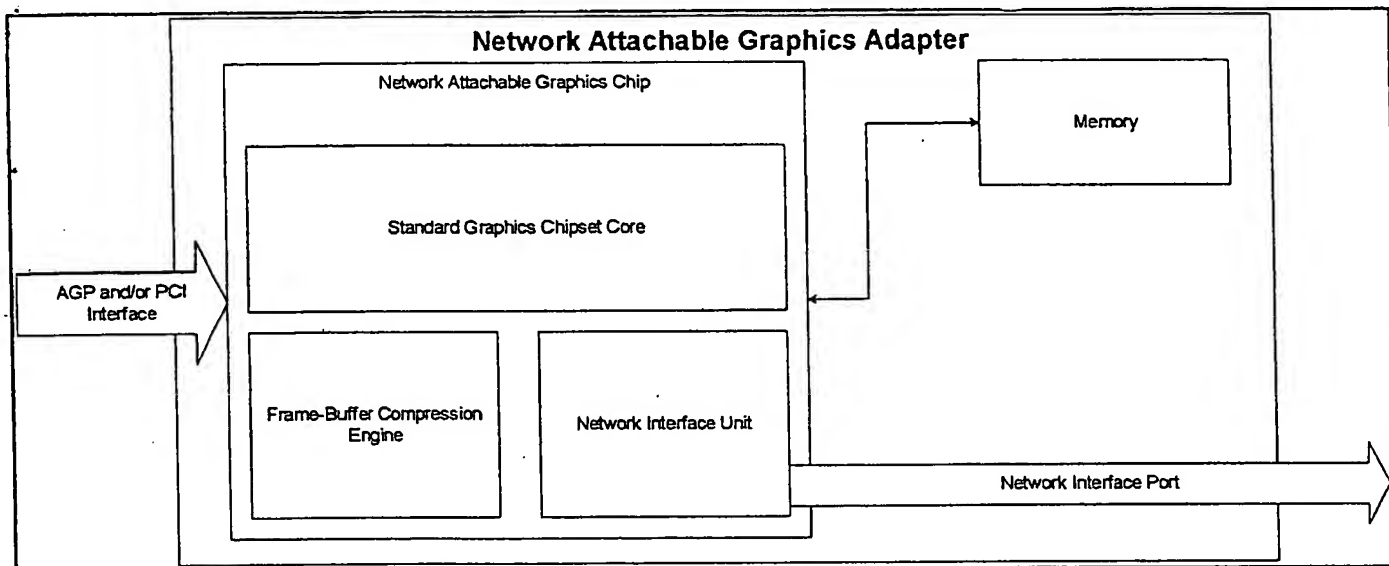


Figure 3 Variation 3

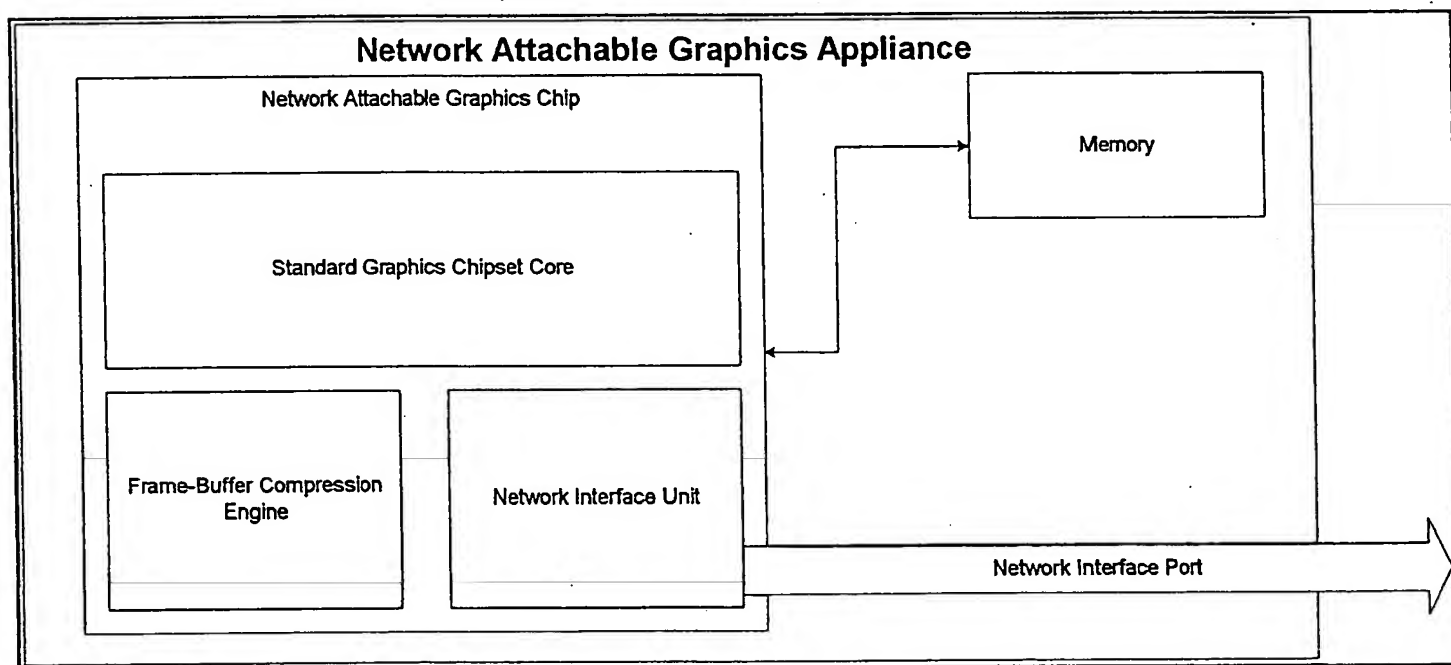


Figure 4 Variation 4